

Exhibit 14

Google AI Blog

The latest from Google Research

Contactless Sleep Sensing in Nest Hub

Tuesday, March 16, 2021

Posted by Michael Dixon, Software Engineer and Reena Singhal Lee, Product Manager, Google Health


People often turn to technology to manage their health and wellbeing, whether it is to record their daily exercise, measure their heart rate, or increasingly, to understand their sleep patterns. Sleep is foundational to a person's everyday wellbeing and can be impacted by (and in turn, have an impact on) other aspects of one's life — mood, energy, diet, productivity, and more.

As part of our ongoing efforts to support people's health and happiness, today we announced Sleep Sensing in the new Nest Hub, which uses radar-based sleep tracking in addition to an algorithm for cough and snore detection. While not intended for medical purposes¹, Sleep Sensing is an opt-in feature that can help users better understand their nighttime wellness using a contactless bedside setup. Here we describe the technologies behind Sleep Sensing and discuss how we leverage on-device signal processing to enable sleep monitoring (comparable to other clinical- and consumer-grade devices) in a way that protects user privacy.

Soli for Sleep Tracking

Sleep Sensing in Nest Hub demonstrates the first wellness application of Soli, a miniature radar sensor that can be used for gesture sensing at various scales, from a finger tap to movements of a person's body. In Pixel 4, Soli powers Motion Sense, enabling touchless interactions with the phone to skip songs, snooze alarms, and silence phone calls. We extended this technology and developed an embedded Soli-based algorithm that could be implemented in Nest Hub for sleep tracking.

Soli consists of a millimeter-wave frequency-modulated continuous wave (FMCW) radar transceiver that emits an ultra-low power radio wave and measures the reflected signal from the scene of interest. The frequency spectrum of the reflected signal contains an aggregate representation of the distance and velocity of objects within the scene. This signal can be processed to isolate a specified range of interest, such as a user's sleeping area, and to detect and characterize a wide range of motions within this region, ranging from large body movements to sub-centimeter respiration.



Soli spectrogram illustrating its ability to detect a wide range of motions, characterized as (a) an empty room (no variation in the reflected signal demonstrated by the black space), (b) large pose changes, (c) brief limb movements, and (d) sub-centimeter chest and torso displacements from respiration while at rest.

In order to make use of this signal for Sleep Sensing, it was necessary to design an algorithm that could determine whether a person is present in the specified sleeping area and, if so, whether the person is asleep or awake. We designed a custom machine-learning (ML) model to efficiently process a continuous stream of 3D radar tensors (summarizing activity over a range of distances, frequencies, and time) and automatically classify each feature into one of three possible states: absent, awake, and asleep.

To train and evaluate the model, we recorded more than a million hours of radar data from thousands of individuals, along with thousands of sleep diaries, reference sensor recordings, and external annotations. We then leveraged the TensorFlow Extended framework to construct a training pipeline to process this data and produce an efficient TensorFlow Lite embedded model. In addition, we created an automatic calibration algorithm that runs during setup to configure the part of the scene on which the classifier will focus. This ensures that the algorithm ignores motion from a person on the other side of the bed or from other areas of the room, such as ceiling fans and swaying curtains.



Wakefulness probability

Search blog ...

Labels

Archive

Feed

Follow @googleai


Give us feedback in our Product Forums.



The custom ML model efficiently processes a continuous stream of 3D radar tensors (summarizing activity over a range of distances, frequencies, and time) to automatically compute probabilities for the likelihood of user presence and wakefulness (awake or asleep).

To validate the accuracy of the algorithm, we compared it to the gold-standard of sleep-wake determination, the polysomnogram sleep study, in a cohort of 33 "healthy sleepers" (those without significant sleep issues, like sleep apnea or insomnia) across a broad age range (19-78 years of age). Sleep studies are typically conducted in clinical and research laboratories in order to collect various body signals (brain waves, muscle activity, respiratory and heart rate measurements, body movement and position, and snoring), which can then be interpreted by trained sleep experts to determine stages of sleep and identify relevant events. To account for variability in how different scorers apply the American Academy of Sleep Medicine's staging and scoring rules, our study used two board-certified sleep technologists to independently annotate each night of sleep and establish a definitive groundtruth.

We compared our Sleep Sensing algorithm's outputs to the corresponding groundtruth sleep and wake labels for every 30-second epoch of time to compute standard performance metrics (e.g., sensitivity and specificity). While not a true head-to-head comparison, this study's results can be compared against previously published studies in similar cohorts with comparable methodologies in order to get a rough estimate of performance. In "Sleep-wake detection with a contactless, bedside radar sleep sensing system", we share the full details of these validation results, demonstrating sleep-wake estimation equivalent to or, in some cases, better than current clinical and consumer sleep tracking devices.




Aggregate performance from previously published accuracies for detection of sleep (sensitivity) and wake (specificity) of a variety of sleep trackers against polysomnography in a variety of different studies, accounting for 3,990 nights in total. While this is not a head-to-head comparison, the performance of Sleep Sensing on Nest Hub in a population of healthy sleepers who simultaneously underwent polysomnography is added to the figure for rough comparison. The size of each circle is a reflection of the number of nights and the inset illustrates the mean±standard deviation for the performance metrics.

Understanding Sleep Quality with Audio Sensing


The Soli-based sleep tracking algorithm described above gives users a convenient and reliable way to see how much sleep they are getting and *when* sleep disruptions occur. However, to understand and improve their sleep, users also need to understand *why* their sleep is disrupted. To assist with this, Nest Hub uses its array of sensors to track common sleep disturbances, such as light level changes or uncomfortable room temperature. In addition to these, respiratory events like coughing and snoring are also frequent sources of disturbance, but people are often unaware of these events.

As with other audio-processing applications like speech or music recognition, coughing and snoring exhibit distinctive temporal patterns in the audio frequency spectrum, and with sufficient data an ML model can be trained to reliably recognize these patterns while simultaneously ignoring a wide variety of background noises, from a humming fan to passing cars. The model uses entirely on-device audio processing with privacy-preserving analysis, with no raw audio data sent to Google's servers. A user can then opt to save the outputs of the processing (sound occurrences, such as the number of coughs and snore minutes) in Google Fit, in order to view personal insights and summaries of their night time wellness over time.





To train the model, we assembled a large, hand-labeled dataset, drawing examples from the publicly available AudioSet research dataset as well as hundreds of thousands of additional real-world audio clips contributed by thousands of individuals.



Log-Mel spectrogram inputs comparing cough (left) and snore (right) audio snippets.

When a user opts in to cough and snore tracking on their bedside Nest Hub, the device first uses its Soli-based sleep algorithms to detect when a user goes to bed. Once it detects that a user has fallen asleep, it then activates its on-device sound sensing model and begins processing audio. The model works by continuously extracting spectrogram-like features from the audio input and feeding them through a convolutional neural network classifier in order to estimate the probability that coughing or snoring is happening at a given instant in time. These estimates are analyzed over the course of the night to produce a report of the overall cough count and snoring duration and highlight exactly when these events occurred.

Conclusion

The new Nest Hub, with its underlying Sleep Sensing features, is a first step in empowering users to understand their nighttime wellness using privacy-preserving radar and audio signals. We continue to research additional ways that ambient sensing and the predictive ability of consumer devices could help people better understand their daily health and wellness in a privacy-preserving way.

Acknowledgements

This work involved collaborative efforts from a multidisciplinary team of software engineers, researchers, clinicians, and cross-functional contributors. Special thanks to D. Shin for his significant contributions to this technology and blogpost, and Dr. Logan Schneider, visiting sleep neurologist affiliated with the Stanford/VA Alzheimer's Center and Stanford Sleep Center, whose clinical expertise and contributions were invaluable to continuously guide this research. In addition to the authors, key contributors to this research from Google Health include Jeffrey Yu, Allen Jiang, Arno Charton, Jake Garrison, Navreet Gill, Sinan Hersek, Yijie Hong, Jonathan Hsu, Andi Janti, Ajay Kannan, Mukil Kesavan, Linda Lei, Kunal Okhandiar, Xiaojun Ping, Jo Schaeffer, Neil Smith, Siddhant Swaroop, Bhavana Koka, Anupam Pathak, Dr. Jim Taylor, and the extended team. Another special thanks to Ken Mixter for his support and contributions to the development and integration of this technology into Nest Hub. Thanks to Mark Malhotra and Shwetak Patel for their ongoing leadership, as well as the Nest, Fit, Soli, and Assistant teams we collaborated with to build and validate Sleep Sensing on Nest Hub.

¹ Not intended to diagnose, cure, mitigate, prevent or treat any disease or condition.

